

METHOD AND APPARATUS FOR VERIFYING THE INTEGRITY OF COMPUTER
NETWORKS AND IMPLEMENTATION OF COUNTER MEASURES

FIELD OF THE INVENTION

This invention relates to a method and apparatus for verifying the integrity of a computer security subsystem for preventing attacks on computer network security systems.

BACKGROUND OF THE INVENTION

Concurrent with the rise in connectivity among diverse computer networks and the corresponding increase in dependence on networked information systems, there has been a dramatic increase in the need for robust security to enforce restrictions on access to and prevent intrusion on secure systems. The topology of the interconnected networks has also grown increasingly complex, and often involves open networks such as the internet that expose secure systems to increased threats of attack. Consequently, no single solution has yet been proposed that addresses all current needs for intrusion detection and response. Instead, a vast assortment of security devices and techniques has evolved and has generally been implemented differently on individual systems. This has resulted in a global security patchwork, inherently susceptible to attack and to individual systems which themselves implement a hodge podge of different security devices and techniques.

Attempts to gain unauthorized access to computer networks capitalize on inherent loopholes in a network's security topology. It is known, for example, that although a secure system connected to the internet may include firewalls and intrusion detection systems to prevent unauthorized access, weaknesses in individual security components are often sought out and successfully exploited. The rapid introduction of new

5 technology exacerbates the problem, creating or exposing additional weaknesses that may not become known until a breach in security has already occurred.

A fundamental weakness shared in common by current intrusion detection and response systems is their "flat" or non-hierarchical implementation. The configuration shown in Fig. 1 is an example of such a typical network implementation on a
10 hypothetical "target network". The network 10 includes a plurality of file servers 14, workstations 16, a network intrusion detection system (IDS) 18, a remote access server 20 and a web server 22. These devices are connected to each other over network backbone 12, and form a local or wide-area network (LAN or WAN). Router 26 is connected directly to an open network such as the internet, 30, and is connected to the devices on network backbone 12 through network firewall 24.

The firewall 24 and the IDS 18 are part of the security system of network 10. Firewall 24 is configurable and serves to control access by hosts on the internet to resources on the network. This protects network 10 from intruders outside the firewall, essentially by filtering them out. IDS 18 scans packets of information transmitted over backbone 12 and is configured to detect specific kinds of transactions that indicate that an intruder is attempting, or already has gained access to the network, 10. In this way, the IDS protects the network from intruders inside as well as outside the firewall. Other devices on network 10 may also contribute to network security, such as remote access
25 server 20 which permits access directly to network 10 from remote computers (not shown), for example over a modem. Remote access server 20 must also implement some security function such as username and password verification to prevent intruders from gaining access to the network and bypassing firewall 24.

In a typical intrusion scenario on a target network connected to the internet, an intruder will first learn as much as possible about the target network from available
30 public information. At this stage, the intruder may do a "whois" lookup, or research DNS tables or public web sites associated with the target. Then, the intruder will engage in a variety of common techniques to scan for information. The intruder may do a "ping"

5 sweep in order to see which machines on the target network are running, or they may
employ various scanning utilities well known in the art such as "rcpinfo", "showmount"
or "snmpwalk" to uncover more detailed information about the target network's
topology. At this stage the intruder has done no harm to the system, but a correctly
configured network IDS should be able, depending on its vantage point on the network,
10 to detect and report surveillance techniques of intruders that follow known patterns of
suspicious activity. These static definitions, known as "intrusion signatures", are
effective only when the intruder takes an action or series of actions that closely follow
the established definitions of suspicious activity. Consequently, if the IDS is not updated,
is disabled or encounters an unknown or new method of attack, it will not respond
properly. However, if steps are not taken at this point in the attack to prevent further
penetration into the target network, the intruder may actually begin to invade the
network, exploiting any security weaknesses (such as the IDS that may not have reacted
earlier to the intruder), and securing a foothold on the network. Once entrenched, the
intruder may be able to modify or disable any device belonging to the target network
including any remaining IDS or firewall.

Methods used by intruders to gain unauthorized access to computer networks
evolve in sophistication in lock step with advances in security technology. It is a typical,
however that successful attacks on network systems often begin by attacking the security
subsystems in place on the target network that are responsible for detecting common
25 intrusion signatures, disabling those systems and destroying evidence of the intrusion.

U.S. Patent No. 5,916,644 to Kurtzberg et al. discloses a method for testing the
integrity of security subsystems wherein a specially configured system connected to
directly a target computer network will systematically test security on the network by
simulating attacks on security devices in order to verify that they are operational.
30 Specifically, the disclosed method randomly simulates an attack on the network. If the
attack is detected, the security subsystems are assumed to be functioning. If not, they
are considered compromised, and an attack may already be underway. This method is

5 an improvement over passive systems that do not check themselves and therefore cannot properly report on their own status when they have been disabled.

10 A major shortcoming of this approach is that these security systems reside on the same networks that they seek to protect and are similarly vulnerable to attack once an intruder has gotten a foothold on the network. In other words, they are not themselves immune to the attacks of intruders. As a result each advance in the prior art is just another new security hurdle on the network to be defeated. In this light, the active scanning approach disclosed in Kurtzberg is not fundamentally different from any other security measure (such as a firewall) in that it is non-hierarchical and depends completely on the vigilance of a human network manager.

Therefore, there exists a need for a self-diagnosing network security system that can protect a target network from both internal and external intruders and that is resistant to attacks perpetrated on the system it has been deployed to protect. Furthermore, there is a need for an active security system that will take measured action against perceived security threats even in the absence of a human network manager.

Summary of the Invention

It is therefore an object of the present invention to provide a network security system for a network of computers that is capable of solving the above mentioned problems in the prior art.

25 It is another object of the present invention to provide a network security system that has a component that can directly monitor multiple network security devices on a network for attack signatures and other suspicious network activity suggesting an attempt to compromise security on that network.

30 It is another object of the present invention to provide a network security system that can dynamically detect new patterns or trends in network activity that suggest an attempt to compromise network security on a single network or on a plurality of otherwise unrelated networks.

5 It is another object of the present invention to provide a network security system that can resist intrusion during an attack on the network.

 It is another object of the present invention to provide a security system providing integrity verification for security devices on a network, and can also reliably verify its own integrity.

10 It is another object of the present invention to provide a security system for a computer network that can take corrective measures after an attack has been detected to prevent an intruder from gaining further access to the network.

 It is another object of the present invention to provide a security system satisfying the above objectives for individual computers connected to an open network.

 According to an example of the present invention, there is provided a network security system to prevent intrusion on a target network having at least one security subsystem local to the target network provided to monitor network traffic and to detect attacks by an intruder on the system. The subsystem is connected via a secure link to a master system that is not otherwise connected to the target system. The master system monitors the subsystem via the secure link and registers information pertaining to the status of the subsystem. If the subsystem detects an attack on the target network, or does not respond to the master system, the master system will take appropriate action, ranging from logging the incident or notifying a network manager to attempting to shut down the network. Accordingly, even attacks that completely disable the subsystem will
25 not prevent the master system from responding as long as the link remains secure.

 According to another example of the present invention, a multi-level hierarchy is implemented making the subsystem subordinate to the master system. In this configuration, commands can only be passed from the master system to the subsystem, ensuring that the integrity of the master system can not be undermined, even by
30 successful attacks on the target network, or on the subsystem itself. Therefore, even a subversion of the subsystem and a compromised link between it and the master system is insufficient to disable the master system.

5 According to another example of the present invention, a pseudo-attack generator associated with the master system is provided that simulates attacks on the target network that should be detected by the subsystem. By comparing the pseudo-attacks made on the target network to the attacks actually detected by the subsystem, the master system can determine whether the integrity of the subsystem has been compromised.
10 Similarly, the subsystem may generate its own pseudo-attacks on other network security components to establish their integrity as well. Therefore it is possible to test comprehensively every security-related device connected to the target network.

In another example of the present invention, the subsystem, and the master system acting through the subsystem, can implement corrective measures to mitigate or thwart suspected intruder attacks on the target network.

BRIEF DESCRIPTION OF THE DRAWINGS:

Fig. 1 is a block diagram showing the overall structure of an example of a network system according to the prior art.

Fig. 2 is a block diagram showing an example of a network incorporating the system of the present invention.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

The preferred embodiments of a network security system according to the present invention will hereinafter be described with reference to the accompanying drawings.

Referring to Fig. 2, a first embodiment of the present invention is shown. Target network 100 is shown having the same basic components as the network of the prior art shown in Fig. 1 with the addition of security subsystem 50, however it should be noted that the actual configuration of the target network is not critical with the exception of at least one security subsystem 50. Each of the security subsystem 50, servers 14,

workstations 16, IDS 18, remote access server 20, web server 22, firewall 24 and router 26 are connected together over network backbone 12. Each of the devices carry out communication over the backbone in accordance with a predetermined communication protocol such as Transmission Control Protocol/Internet Protocol (TCP/IP).

Target network 100 is connected through firewall 24 and router 26 to the internet 30 as well as through remote access server 20 which may also be selectively connected to the internet 30 through remote user 21. These two potential points of contact with an open network, in this case the internet, exposes target network 100 to the threat of intrusion from any host with access to the internet such as internet user 31. In addition to threats from the outside, those with direct access to the resources of target network 100, such as those using one of the workstations 16, also pose an intrusion threat. If an intruder were to gain access to one of the critical security-related devices such as the IDS 18 or the firewall 24 or any trusted computer from within or outside the target network 100, security on the network could be compromised.

In the present invention, security subsystem 50 is connected to network backbone 12 and linked to each of the network's devices by a secure link 52. Such a secure link may be established through an encrypted communication protocol such as Secure Sockets Layer (SSL). This ensures that communication between the security subsystem 50 and the other components of the target network cannot be intercepted by an intruder. A similar secure link 54 is established as a virtual private network (VPN) tunnel between the security subsystem 50 and a master system 60 connected to a remote network 110. Although the remote network is shown having its own firewalls 62, servers 66, and router 68, the ultimate configuration of remote network 110 is not critical beyond secure link 54 connecting security subsystem 50 and master system 60. However, secure links 55 may be established between a device such as a network scanner 63 and a router 26 or remote user 21 on network 100. Secure link 54 ensures that communication between the two networks cannot be intercepted by an intruder. Therefore, there should be no

5 other direct connection between target network 100 and remote network 110 except over a secure link.

Sub D3 → ~~Preferably, the security system defined herein is embedded as a software package and implemented on computers comprising at least a master system and the security subsystem.~~

10 During operation, security subsystem 50 monitors the activities of the devices of the target network 100. Particularly, the critical security-related functions of IDS 18 and firewall 24 are tested. The particular method employed by security subsystem 50 in testing these devices is not critical, however the above mentioned approach employing simulated attacks on the components would be suitable.

UPON TESTING OF DEVICES
20 Upon testing the devices, if the integrity of a device on target network 100 cannot be verified, security subsystem 50 reacts. For example, if IDS 18 has been identified by the subsystem as not reacting properly to attacks on it originating from the internet, appropriate countermeasures could include cutting off or restricting access to the network at firewall 24 or stop at application level. If instead, the firewall is determined not to be functioning, appropriate action might include disabling access to any servers 14 holding sensitive data. In one possible configuration of the present invention, security subsystem 50 reports network device status to master system 60 which processes the information, and decides on further action. In an alternate configuration, security subsystem 50 is responsible for implementing countermeasures directly. In both cases,
25 however, the results of every test are passed to master system 60 where they are stored for analysis.

30 The system of the present invention can also help thwart ongoing attacks and is uniquely suited to do so. In another preferred embodiment of the present invention, master system 60 hierarchically supercedes security subsystem 50. As such, the activities of security subsystem 50 are defined as a child process of master system 60 and are subordinate thereto. Although information preferably flows both ways between master

5 system 60 and security subsystem 50 in this embodiment, the master system in this embodiment does not take direction from the subsystem.

As noted in the discussion of the prior art, non-hierarchical security systems are connected directly to a target network and are inherently susceptible to attacks on that network. This is in contrast to the present embodiment wherein, even if completely
10 subverted during an attack on target system 100, security subsystem 50 would not result in a takeover of master system 60. The benefit of this configuration is that the master system would still be able to carry out its function. For example, if master system 60 is configured to sound an alarm when security subsystem 50 no longer responds to it, there would be no way, in this embodiment, for intruders on target network 100 to remotely shut down master system 60 because the master system will not respond to any
15 instructions issued from a subordinate system. Although master system 60 may lose control of the target network, it is not in danger of being taken over by it. Additionally, if the link 54 between master system 60 and security subsystem 50 is severed or compromised, instructions may be routable instead through secure links 55.

In yet another embodiment of the present invention, remote network 110 is connected through router 70 to an open network such as the Internet. This enables master system 60 to send random pseudo-attacks to target network 100. The pseudo-attacks may mimic any of the actual attack signatures known by the master system to be detectable by the target network. If the expected reply is not received by the master
20 system, an early indication of an intruder attack on the target network is indicated.

As set forth hereinabove, according to the present invention, it is possible to provide a method and apparatus for verifying the integrity of computers and computer networks that is independent of the network or computer being tested. In addition, by detecting early signs of intruder activity on a network, the present invention increases
25 the likelihood that intruder attacks can be thwarted before they succeed.

When implemented on an individual computer, such as a single workstation 16 connected to an open network such as internet 30, the present invention functions

5

10

Many different embodiments of the present invention may be constructed without departing from the spirit and scope of the invention. It should be understood that the present invention is not limited to the specific embodiments described in this specification. To the contrary, the present invention is intended to cover various modifications and equivalent arrangements included within the spirit and the scope of the claims.